In preparation manuscript No.
(will be inserted by the editor)

Anisotropic Diffusion Filtering of Images on Curved Surfaces

Emma Naden · Thomas März · Colin B. Macdonald

arXiv:1403.2131v2 [math.NA] 5 Apr 2014

Received: April 8, 2014/ Accepted: date

Abstract We demonstrate a method for filtering images defined on curved surfaces embedded in 3D. Applications are noise removal and the creation of artistic effects. Our approach relies on in-surface diffusion: we formulate Weickert's edge/coherence enhancing diffusion models in a surface-intrinsic way. These diffusion processes are anisotropic and the equations depend nonlinearly on the data. The surface-intrinsic equations are dealt with the closest point method, a technique for solving partial differential equations (PDEs) on general surfaces. The resulting algorithm has a very simple structure: we merely alternate a time step of a 3D analog of the in-surface PDE in a narrow 3D band containing the surface with a reconstruction of the surface function. Surfaces are represented by a closest point function. This representation is flexible and the method can treat very general surfaces. Experimental results include image filtering on smooth surfaces, open surfaces, and general triangulated surfaces.

Keywords partial differential equations · denoising · feature enhancement · surfaces · closest point method

Emma Naden
Leidos Health
705 E. Main St.
Westfield, IN 46074, USA
E-mail: emma.naden@oxon.org

Thomas März
Mathematical Institute
University of Oxford
Oxford OX2 6GG, UK
E-mail: maerz@maths.ox.ac.uk

Colin B. Macdonald
Mathematical Institute
University of Oxford
Oxford OX2 6GG, UK
E-mail: macdonald@maths.ox.ac.uk

1 Introduction

Image processing is an active field of mathematical research, and models based on partial differential equations (PDEs) have been used successfully for various image processing tasks. Examples include inpainting [6, 12, 27], segmentation [13, 19, 40], smoothing and denoising [31, 38, 11, 2].

Denoising on general curved surfaces [23], the corresponding scale space analysis [21], and related image processing problems have seen some interest [36, 8]. Of particular relevance for our voxel-based technique is that image processing on curved surfaces can occur even when data is acquired in three-dimensional volumes. For example, Lin and Seales [24] propose CT scanning of scrolls in order to non-destructively read text written on rolled up documents. Other applications of denoising on surfaces include digital image-based elasto-tomography (DIET), a technique for non-invasive breast cancer screening [9], texture processing [14, 4], and surface fairing where the surface data is itself a height-field perturbation relative to a reference surface [15].

In the cited articles a common theme is the use of PDE models. The topic of solving PDEs on curved surfaces is an important area of research with many additional applications in physics and biology (for example to model pattern formation via reaction-diffusion on animal coats [29, 34] or the diffusion of chemicals on the surface of a cell [35, 30, 17]).

The numerical treatment of PDEs on surfaces requires a suitable representation of the geometry. The choice of the representation is fundamental to the complexity of the algorithm. Parametrizations are commonly used, however they can introduce artificial distortions and singularities into the models even if the surface

geometry is as simple as that of a sphere. Moreover, joining up multiple patches is typically necessary if the geometry is more complex. On triangular meshes approximating a surface, discretizations can be obtained using finite elements, e.g., [16]. Embedding techniques based on implicit representations such as level sets (e.g., [5]) or the closest point representation are further alternatives and tend to be quite flexible. However with level sets it is not obvious how to treat open surfaces and one has to introduce artificial boundary conditions at the boundary of the 3D neighborhood in the embedding space.

In this work we use the closest point method [33] for solving the diffusion models on surfaces. This technique keeps the resulting evolution as simple as possible by alternating between two straightforward steps:

1. A time step of the analogous three-dimensional PDE model using standard finite difference methods.
2. An interpolation step, which reconstructs the surface function and makes the 3D calculation consistent with the surface problem.

The surface geometry is encoded in a closest point function [33,28] which maps every off-surface point to that surface point which is closest in Euclidean distance. The closest point function is used in the reconstruction step only and therefore this technique does not require modification of the model via a parameterization nor does it require the surface to be closed or orientable. The closest point method has already been used successfully in image segmentation [37], visual effects [3] and to perform Perona–Malik edge-stopping diffusion [7].

The use of diffusion processes to smooth noisy images dates back to at least the 1980s, e.g., Koenderink [22] notes that one could perform smoothing by solving the heat equation over the image domain. Unfortunately, the linearity of this equation implies a uniform blur of the entire image without any consideration of the structure of the image. Consequently, important high-frequency components due to edges are blurred as much as undesired high-frequency components due to noise and the structure of the image is lost. Perona and Malik [31] put forth an elegant and simple modification of the heat equation to address this issue. In order to avoid blurring edges, they vary the rate of diffusivity according to the local gradient of the image, stopping the diffusion process at edges (i.e., where the magnitude of the gradient is large). Regularizations of the Perona–Malik model and other variations of the diffusivity (and also relations to geometric evolutions such as mean curvature motion) have been studied for example in [11,2]. Although there is an improvement over linear diffusion, Perona and Malik's method retains a fundamental flaw: because the diffusion coefficient drops to zero when the gradient is high, their technique cannot remove noise along edges. This motivates the use of an anisotropic diffusion filter as suggested by Weickert [38] in which one alters not only the rate of diffusion near an edge but also its direction so that diffusion is performed along rather than across the edge. In order to accomplish this, the scalar diffusion coefficient found in linear and nonlinear Perona–Malik diffusion must be replaced with a diffusion tensor. Apart from edge-enhancing denoising, Weickert's diffusion model [38] has a second application called coherence-enhancing diffusion which can be used to create artistic effects. The contribution of this paper is to extend Weickert's model to curved surfaces and implement anisotropic diffusion on surfaces using the closest point method. Anisotropic diffusion to smooth surface meshes as well as functions defined on surface meshes has been considered in [4]. Their approach differs from ours in two main ways. Firstly, the design of their diffusion tensor is based on the directions of principal curvature while ours is based on eigenvectors of a surface-intrinsic structure tensor. Secondly, they use surface finite elements to discretize the diffusion equation while we use the closest point method.

The rest of this paper unfolds as follows. Section 2 reviews the concepts of anisotropic diffusion filtering closely following Weickert [38]. In Section 3 we introduce a mathematical model for surface-intrinsic anisotropic diffusion. In particular, we define the structure tensor for images on surfaces which allows us to estimate visual edges. Section 4 introduces the closest point method for solving diffusion PDEs posed on surfaces and describes our numerical schemes. The presentation of results and a discussion follows in Section 5.

2 Anisotropic Diffusion in Image Processing

A *digital image* is a quantitative representation of the information (color and intensity values) encoded in an image. Typically, the domain of a digital image is a finite and discrete set of points called *pixels*. We view images as a functions defined on a rectangular domain in $\mathbb{R}^2$. Digital images are then discretizations of such functions obtained from sampling on a Cartesian grid. We will later extend this to images defined on two-dimensional surfaces.

When we view images as a functions, many image processing tasks can be described by time-dependent PDE processes such as diffusion. The resulting image is then the solution to the PDE. Color images are represented by vector-valued functions, for example, in RGB color space we have a triple $u(x) = [R, G, B]$ at each point x in the domain, where $R, G, B \in [0, 1]$ represent the intensities of red, green, and blue respectively.

When working with color images, we apply the diffusion filter to each of the color components individually. Unless otherwise noted, throughout the remainder of this paper we assume $u \in [0,1]$, where zero represents black (no color intensity) and one represents white (full color intensity).

A so-called *noisy image* is an image that differs from the real object depicted (the *ground truth*), by a discoloration of pixels in the digital image. Noise can occur from a variety of sources depending on the process used to generate or transmit the digital image. For obvious reasons, it is desirable to remove the noise in a visually plausible manner.

2.1 Anisotropic Diffusion

In PDE-based diffusion filtering we typically process the image by solving a PDE of the form

$$\partial_t u = \operatorname{div}(\mathbf{j}) \qquad \text{in } \Omega, t > 0, \tag{1a}$$

$$u|_{t=0} = u^0, \tag{1b}$$

$$\mathbf{j}^{\mathrm{T}}\boldsymbol{\nu} = 0 \qquad \text{on } \partial\Omega, t > 0\,. \tag{1c}$$

Here, u^0 is the initially given image, Ω denotes the rectangular image domain, $\boldsymbol{\nu}$ is the exterior boundary normal, and $\mathbf{j}$ is the flux vector. For isotropic filters the flux vector takes the form

$$\mathbf{j} = \kappa \nabla u\,.$$

If we choose $\kappa = 1$ here we have Gaussian diffusion. If we choose the Perona–Malik function $\kappa = g(|\nabla u|^2)$ we would have Perona–Malik[1] edge-stopping diffusion.

Depending on the choice of κ, isotropic filters can blur an edge (e.g., Gaussian diffusion), displace an edge, or just stop the diffusion there (e.g., Perona–Malik). However isotropic filters cannot smooth noise along an edge. This is because the flux is parallel to the gradient, which is perpendicular to edges.

An anisotropic flux vector could take the form

$$\mathbf{j} = \kappa_1 \nabla u + \kappa_2 \nabla u^{\perp}$$

where κ_1 and κ_2 are chosen based on the local structure of the image. For example, near an edge, it would be desirable to prevent smoothing across the edge (κ_1 small) and instead smooth only along the edge (κ_2 large).

Weickert's approach [38] to implement this idea was to design a diffusion tensor

$$\mathsf{G}[u] = \kappa_1 \boldsymbol{\omega}^{\perp}\boldsymbol{\omega}^{\perp\mathrm{T}} + \kappa_2 \boldsymbol{\omega}\boldsymbol{\omega}^{\mathrm{T}}, \tag{2}$$

[1] In their paper, Perona and Malik referred to their method as anisotropic although we consider it to be isotropic in the sense that the flux is always parallel to the gradient.

where $\boldsymbol{\omega} = \boldsymbol{\omega}[u]$ is a normalized "edge vector" that approximates the visual edge in the image u in a small neighborhood about the point x. The corresponding flux is given by

$$\mathbf{j} = \mathsf{G}[u]\nabla u \;= \kappa_1 \left(\boldsymbol{\omega}^{\perp\mathrm{T}}\nabla u\right)\, \boldsymbol{\omega}^{\perp} + \kappa_2 \left(\boldsymbol{\omega}^{\mathrm{T}}\nabla u\right)\, \boldsymbol{\omega} \tag{3}$$

and is directed by the visual edge $\boldsymbol{\omega}$ rather than the gradient direction. The tensor $\mathsf{G}[u] \in \mathbb{R}^{2\times 2}$ is a nonlinear function of u and varies spatially with x. The diffusion is anisotropic since G is not simply a scalar multiple of the identity matrix.

The visual edge vector $\boldsymbol{\omega}$ needed to construct G is found by *structure tensor analysis* [38]. The structure tensor and the different choices for κ_1, κ_2 in (2)—which lead to edge- and coherence-enhancing diffusion—are discussed briefly in the following.

2.2 The Structure Tensor

The structure tensor provides information about orientations and coherent structures in an image [38,10,1]. In its simplest form, the structure tensor is the matrix

$$\mathsf{J}_{\sigma,0}[u] := \nabla u_\sigma \nabla u_\sigma^{\mathrm{T}},$$

where—in order to reduce the effect of noise or irrelevant small scale features—the image u has been smoothed with a heat-kernel

$$u_\sigma := K_\sigma * u, \qquad K_\sigma(x) = \frac{e^{\left(-\frac{x^2}{4\sigma}\right)}}{\sqrt{4\pi\sigma}}.$$

In contrast to [38] we prefer the use of heat kernels K_σ over Gaussians as this will extend in a straightforward manner to surfaces (see Section 3), and has the same effect in practice. $\mathsf{J}_{\sigma,0}[u]$ is the initial form of the structure tensor; coherent structures over a bigger neighborhood around the point x are often accounted for by averaging tensor-valued orientation information given in $\mathsf{J}_{\sigma,0}$. This is done by component-wise convolution of $\mathsf{J}_{\sigma,0}$ with a second heat-kernel K_ρ to define the structure tensor

$$\mathsf{J}_{\sigma,\rho}[u] := K_\rho * \mathsf{J}_{\sigma,0}[u] = K_\rho * \left(\nabla u_\sigma \nabla u_\sigma^{\mathrm{T}}\right). \tag{4}$$

By construction, $\mathsf{J}_{\sigma,\rho}[u]$ is symmetric positive semidefinite, and has therefor a spectral decomposition

$$\mathsf{J}_{\sigma,\rho}[u] = \mu_1 \boldsymbol{\omega}^{\perp}\boldsymbol{\omega}^{\perp\mathrm{T}} + \mu_2 \boldsymbol{\omega}\boldsymbol{\omega}^{\mathrm{T}}, \quad \mu_1 \geq \mu_2 \geq 0. \tag{5}$$

This decomposition is unique if $\mu_1 \neq \mu_2$ and yields the orientations parallel $\boldsymbol{\omega}$ and perpendicular $\boldsymbol{\omega}^{\perp}$ to the visual edge. (Note that if $\rho = 0$, one obtains $\mu_2 = 0$

and $\boldsymbol{\omega} || \nabla u_\sigma^\perp$.) The eigenvalues give information about the image structure near x, specifically (cf. [38])

$$\mu_1 = \mu_2 \quad \Longrightarrow \quad x \text{ in homogeneous region}, \tag{6a}$$
$$\mu_1 \gg \mu_2 = 0 \quad \Longrightarrow \quad x \text{ near straight edge}, \tag{6b}$$
$$\mu_1 \geq \mu_2 \gg 0 \quad \Longrightarrow \quad x \text{ near corner}. \tag{6c}$$

The so-called *coherence* $c := \mu_1 - \mu_2 \geq 0$ measures how pronounced the edge is. It also gives a measure for the local contrast and tells us if the decomposition (5) is well- or ill-conditioned.

Now, we consider the choices of κ_1, κ_2 in (2) to obtain the diffusion tensor $\mathsf{G}[u]$. The eigenvectors $\boldsymbol{\omega}^\perp$, $\boldsymbol{\omega}$ of $\mathsf{G}[u]$ are exactly those of the structure tensor $\mathsf{J}_{\sigma,\rho}[u]$ from (5), while we replace the eigenvalues with κ_1, κ_2, depending on the coherence c.

2.3 Edge-Enhancing Diffusion

In edge-enhancing diffusion, the goal is to smooth noise while keeping or enhancing edges. According to Weickert [38], one works with the initial structure tensor $\mathsf{J}_{\sigma,0}[u]$. The appropriate choice of κ_1 and κ_2 is based on the following observation: in homogeneous regions $|\nabla u_\sigma|^2 = \mu_1 \approx \mu_2 = 0$, whereas near an edge $\mu_1 \gg \mu_2 = 0$. Thus, κ_1 (the diffusivity across edges) is chosen to be a decreasing function of μ_1, while κ_2 (the diffusivity along edges) is set equal to one:

$$\kappa_1 = g(\mu_1) \quad \text{and} \quad \kappa_2 = 1. \tag{7}$$

Here g is a scalar-valued function with $g(0) = 1$ and $\lim_{s\to\infty} g(s) = 0$. We choose the Perona–Malik diffusivity function

$$g(s^2) = \frac{1}{1 + \frac{s^2}{\lambda^2}}, \tag{8}$$

where $\lambda \in \mathbb{R}$ in is tunable parameter that determines the filter's sensitivity to edges. For a different choice of g see [38]. In regions of very little contrast, we have $\kappa_1 \approx 1$ and the diffusion will behave locally like Gaussian diffusion since $\mathsf{G}[u] \approx I$. Near an edge, we have $\kappa_1 \approx 0$ which implies $\mathsf{G}[u] \approx \boldsymbol{\omega}\boldsymbol{\omega}^{\mathrm{T}}$, and diffusion occurs mainly along the edge.

The edge-enhancing set-up of $\mathsf{G}[u]$ can also be based on $\mathsf{J}_{\sigma,\rho}[u]$. But with $\rho > 0$, $\kappa_1 = g(c)$ is a function of the coherence $c := \mu_1 - \mu_2$. Because the decomposition (5) is ill-conditioned if $c \approx 0$, we force $\mathsf{G}[u] \approx \mathsf{I}$ in that case. This is also compatible with the case $\rho = 0$, where $c = \mu_1$ because $\mu_2 = 0$. An example of edge-enhancing diffusion filtering is shown in Figure 1.

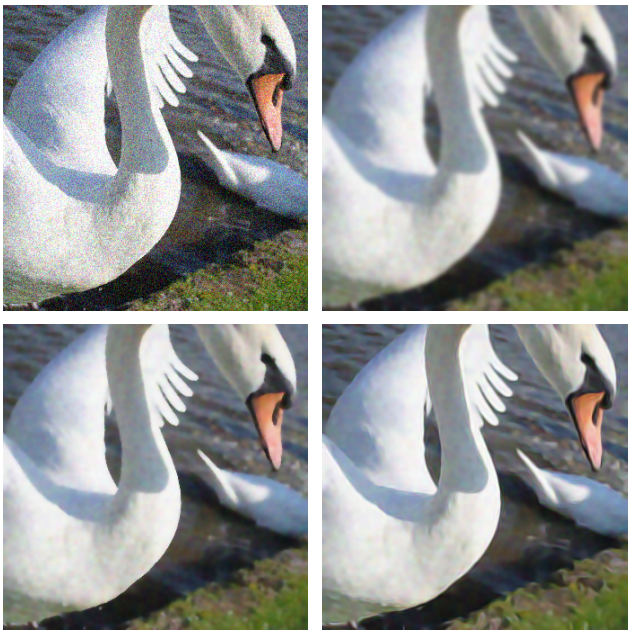

Fig. 1: The effect of different diffusion filters. *Top Left:* Noisy image (photograph by courtesy of Harry Biddle). *Top Right:* Gaussian diffusion. *Bottom Left:* Perona–Malik diffusion; edge-sensitivity parameter $\lambda_{\mathrm{rel}} = 0.3$. *Bottom Right:* Edge-enhancing diffusion; $\sigma = 0.75$, $\rho = 2.5$, $\lambda_{\mathrm{rel}} = 10^{-2}$. Stop time in all cases $T = 2.5$.

2.4 Coherence-Enhancing Diffusion

The second type of anisotropic diffusion which we consider is coherence-enhancing diffusion which aims to preferentially smooth in the direction of largest coherence. Weickert [38] suggests the following choice for κ_1 and κ_2 in (2):

$$\kappa_1 = \alpha, \qquad \kappa_2 = \alpha + (1-\alpha)\exp\left(\frac{-B^2}{(\mu_1 - \mu_2)^2}\right), \tag{9}$$

where $0 \leq \alpha \ll 1$ and $B > 0$ are constants. κ_2 is a function of the coherence $c = \mu_1 - \mu_2$ where μ_1 and μ_2 are the eigenvalues of $\mathsf{J}_{\sigma,\rho}[u]$. Diffusion along the edge (κ_2 large) is performed when the coherence is large compared to B. Where the coherence is small, i.e., $\kappa_2 \approx \kappa_1 = \alpha$, only a small amount of Gaussian diffusion is performed.

One application of coherence-enhancing diffusion lies in its ability to create stylized images (cf. [38]) which might be desired for use in comic books or video games. Coherence-enhancing diffusion creates a natural looking texture because of smoothing according to the inherent structure of the image. As the stop time increases, the image becomes less photorealistic and more stylized; this is shown in Figure 2.

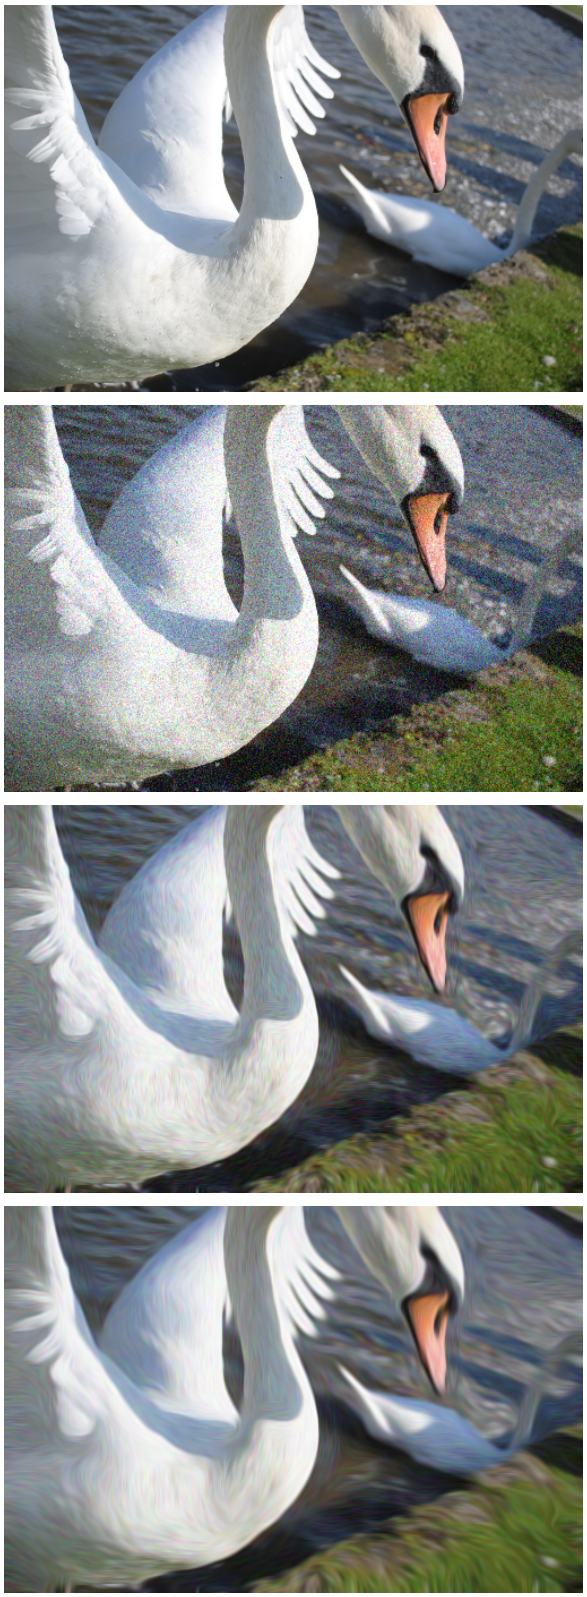

Fig. 2: Stylized texture created by coherence-enhancing diffusion. From top to bottom: *1.* Original photograph. *2.* Noise added for texture creation. *3.* Stop time $T = 2.5$. *4.* Stop time $T = 7.5$. The parameters are $\sigma = 0.75$, $\rho = 25$, $\alpha = 10^{-3}$, and $B_{\mathrm{rel}} = 10^{-10}$.

2.5 Color Images

If the given image u^0 is a color image the diffusion is applied separately to each channel u_k

$$\partial_t u_k = \operatorname{div}\left(\mathsf{G}[u]\nabla u_k\right) \ . \tag{10}$$

In order to minimize spurious colors, we use the same diffusion tensor $\mathsf{G}[u]$ for all channels which is found (as in the scalar case) from a common structure tensor. We compute the common structure tensor following [39] as

$$\mathsf{J}_{\sigma,\rho}[u] := \frac{1}{n}\sum_{k=1}^{n} \mathsf{J}_{\sigma,\rho}[u_k], \tag{11}$$

i.e., the arithmetic mean of all channel structure tensors $\mathsf{J}_{\sigma,\rho}[u_k]$. Depending on the color space some channels might be more important to structure information than others. In this case a general weighted average can be used which puts larger weights on the important channels.

3 Surface-Intrinsic Anisotropic Diffusion

In this section we formulate Weickert's diffusion models for images defined on surfaces. We consider two-dimensional surfaces $\mathcal{S}$ embedded in $\mathbb{R}^3$ which are smooth and orientable with a normal field $\mathbf{n}$. Images are now functions $u : \mathcal{S} \to \mathbb{R}$ defined on $\mathcal{S}$. The diffusion PDE to process the image takes the form

$$\partial_t u = \operatorname{div}_{\mathcal{S}}\left(\mathsf{G}[u]\nabla_{\mathcal{S}} u\right) \qquad \text{in } \mathcal{S}, t > 0, \tag{12a}$$

$$u|_{t=0} = u^0, \tag{12b}$$

$$\left(\mathsf{G}[u]\nabla_{\mathcal{S}} u\right)^{\mathrm{T}}\boldsymbol{\nu} = 0 \qquad \text{on } \partial\mathcal{S}, t > 0\ . \tag{12c}$$

The boundary condition (12c) is relevant only in cases where the surface is open. In this case $\boldsymbol{\nu}$ is an exterior co-normal, i.e., it is perpendicular to the boundary curve but tangent to the surface $\boldsymbol{\nu}^{\mathrm{T}}\mathbf{n} = 0$.

The differential operators in PDE (12a) are intrinsic to the surface and are related to the standard operators as follows (see e.g., [28])

$$\nabla_{\mathcal{S}} u = \nabla u - (\mathbf{n}^{\mathrm{T}}\nabla u)\mathbf{n}, \tag{13a}$$

$$\operatorname{div}_{\mathcal{S}}(\mathbf{j}) = \operatorname{div}(\mathbf{j}) - \mathbf{n}^{\mathrm{T}} D\mathbf{j}\ \mathbf{n}, \tag{13b}$$

where $D\mathbf{j}$ is the Jacobian matrix of the flux $\mathbf{j}$. Note that, with this form, the gradient $\nabla_{\mathcal{S}} u$ and the flux $\mathbf{j}$ are both three-component vectors while $\mathsf{G}[u]$ is a symmetric 3×3-matrix.

While not strictly required, it may be desirable that the PDE (12a) represent a surface conservation law (for example, to guarantee that at least the continuous problem conserves overall gray levels). This means

that the flux vector $\mathbf{j} = \mathsf{G}[u]\nabla_{\mathcal{S}} u \in \mathbb{R}^3$ should always be tangent to the surface $\mathcal{S}$, i.e., $\mathbf{j}^{\mathrm{T}}\mathbf{n} = 0$. This is because the divergence theorem holds only for the tangential part $\mathbf{j}_{\tan}$:

$$\int_{\partial\Omega} \mathbf{j}_{\tan}^{\mathrm{T}} \boldsymbol{\nu} d\mathcal{L} = \int_{\Omega} \operatorname{div}_{\mathcal{S}}(\mathbf{j}_{\tan}) d\mathcal{A}, \tag{14}$$

where $\Omega \subset \mathcal{S}$, $d\mathcal{L}$ and $d\mathcal{A}$ are the infinitesimal length and area measures. Since the co-normal is tangent to the surface we have $\mathbf{j}^{\mathrm{T}}\boldsymbol{\nu} = \mathbf{j}_{\tan}^{\mathrm{T}}\boldsymbol{\nu}$, but (14) does not hold when $\mathbf{j}_{\tan}$ is replaced with $\mathbf{j}$ (unless they are equal). Thus only a tangential flux $\mathbf{j} = \mathbf{j}_{\tan}$ will turn PDE (12a) into a surface conservation law.

The design of a diffusion tensor $\mathsf{G}[u]$ which obeys Weickert's model and maps vectors to the tangent space of $\mathcal{S}$ to give a tangential flux is the topic of the next section.

3.1 The Visual Edge Vector on a Surface

We begin with finding visual-edge vectors for functions u defined on a surface. We leverage the approach of [1] to find a surface intrinsic form of the structure tensor. Let $\boldsymbol{\omega}$ again denote the edge vector and let $\Omega(x) \subset \mathcal{S}$ denote a neighborhood of the point $x \in \mathcal{S}$, then we require

$$(\boldsymbol{\omega}^{\mathrm{T}}\nabla_{\mathcal{S}} u(y))^2 \approx 0 \quad \forall\, y \in \Omega(x), \tag{15a}$$
$$|\boldsymbol{\omega}|^2 = 1, \tag{15b}$$
$$\boldsymbol{\omega}^{\mathrm{T}}\mathbf{n}(x) = 0. \tag{15c}$$

Requirement (15a) says that we want a vector $\boldsymbol{\omega}$ which is almost perpendicular to gradients in a neighborhood of x. Condition (15b) excludes trivial solutions $\boldsymbol{\omega} = 0$. Finally, $\boldsymbol{\omega}$ must be tangent to the surface at x which is expressed in (15c).

As in [1], we formulate (15a) as least squares problem. Let $K_\rho(x, y)$ denote the surface heat kernel on $\mathcal{S}$ with ρ proportional to the diameter of $\Omega(x)$, then we have

$$\boldsymbol{\omega} = \arg\min_{\mathbf{v}} \int_{\mathcal{S}} K_\rho(x, y)(\mathbf{v}^{\mathrm{T}}\nabla_{\mathcal{S}} u(y))^2 \, d\mathcal{A}(y), \tag{16}$$

where the minimization happens over all vectors $\mathbf{v}$ satisfying (15b) and (15c). With the 3×3 structure tensor

$$\mathsf{J}_{0,\rho}[u](x) := \int_{\mathcal{S}} K_\rho(x, y)\nabla_{\mathcal{S}} u(y)\nabla_{\mathcal{S}} u(y)^{\mathrm{T}} \, d\mathcal{A}(y) \tag{17}$$

we can rewrite (16) as

$$\boldsymbol{\omega} = \arg\min_{\mathbf{v}} \mathbf{v}^{\mathrm{T}} \, \mathsf{J}_{0,\rho}[u](x) \, \mathbf{v}. \tag{18}$$

If we consider u_σ, i.e., u smoothed with the surface heat kernel $K_\sigma(x, y)$, in place of u in (15a) we will get

$$\mathsf{J}_{\sigma,\rho}[u](x) := \int_{\mathcal{S}} K_\rho(x, y)\nabla_{\mathcal{S}} u_\sigma(y)\nabla_{\mathcal{S}} u_\sigma(y)^{\mathrm{T}} \, d\mathcal{A}(y) \tag{19}$$

as the structure tensor. The new problem is now

$$\boldsymbol{\omega} = \arg\min_{\mathbf{v}\in\mathbb{R}^3} \mathbf{v}^{\mathrm{T}} \, \mathsf{J}_{\sigma,\rho}[u](x) \, \mathbf{v}, \tag{20a}$$
$$|\mathbf{v}|^2 = 1, \tag{20b}$$
$$\mathbf{v}^{\mathrm{T}}\mathbf{n}(x) = 0, \tag{20c}$$

and can be solved with Lagrange multipliers which leads to the following condition

$$\mathsf{J}\,\mathbf{v} - \lambda\mathbf{v} - \mu\mathbf{n} = 0 \tag{21}$$

subject to the side conditions (20b) and (20c). Here we use J and $\mathbf{n}$ as short hands for $\mathsf{J}_{\sigma,\rho}[u](x)$ and $\mathbf{n}(x)$.

In order to reduce the system, we use an orthonormal basis of the tangent space $\mathcal{T}_x\mathcal{S} = \operatorname{span}\{\mathbf{q}_1(x), \mathbf{q}_2(x)\}$. For $\tilde{\mathbf{v}} \in \mathbb{R}^2$ we set

$$\mathbf{v} = \mathsf{Q}\tilde{\mathbf{v}} \quad \text{with} \quad \mathsf{Q} := [\mathbf{q}_1|\mathbf{q}_2] \in \mathbb{R}^{3\times 2}, \tag{22}$$

then (20c) is satisfied and condition (20b) implies $|\tilde{\mathbf{v}}|^2 = 1$. Moreover, (21) turns into

$$\mathsf{J}\,\mathsf{Q}\tilde{\mathbf{v}} - \lambda\mathsf{Q}\tilde{\mathbf{v}} - \mu\mathbf{n} = 0. \tag{23}$$

Finally, left multiplication with Q^{T} reduces (23) to a 2×2 eigenvalue problem

$$\mathsf{Q}^{\mathrm{T}}\mathsf{J}\,\mathsf{Q}\tilde{\mathbf{v}} - \lambda\tilde{\mathbf{v}} = 0. \tag{24}$$

The 2×2 tensor $\tilde{\mathsf{J}}$ given by the tensor contraction

$$\tilde{\mathsf{J}} := \mathsf{Q}^{\mathrm{T}}\mathsf{J}\,\mathsf{Q} \tag{25}$$

can be seen as the surface-intrinsic structure tensor. Now, the spectral decomposition of $\tilde{\mathsf{J}}$

$$\tilde{\mathsf{J}} = \mu_1\tilde{\boldsymbol{\omega}}^{\perp}\tilde{\boldsymbol{\omega}}^{\perp\mathrm{T}} + \mu_2\tilde{\boldsymbol{\omega}}\tilde{\boldsymbol{\omega}}^{\mathrm{T}}, \quad \mu_1 > \mu_2 \tag{26}$$

yields the desired solution: the eigenvector $\tilde{\boldsymbol{\omega}} \in \mathbb{R}^2$ with respect to the minimal eigenvalue μ_2 represents the visual edge in tangent space coordinates, hence in embedding space coordinates we have

$$\boldsymbol{\omega} = \mathsf{Q}\tilde{\boldsymbol{\omega}} \quad \text{and} \quad \boldsymbol{\omega}^{\perp} = \mathsf{Q}\tilde{\boldsymbol{\omega}}^{\perp}. \tag{27}$$

The surface-intrinsic 2×2 diffusion tensor $\tilde{\mathsf{G}}$ is constructed by

$$\tilde{\mathsf{G}} := \kappa_1\tilde{\boldsymbol{\omega}}^{\perp}\tilde{\boldsymbol{\omega}}^{\perp\mathrm{T}} + \kappa_2\tilde{\boldsymbol{\omega}}\tilde{\boldsymbol{\omega}}^{\mathrm{T}}, \tag{28}$$

where we choose κ_1, κ_2 as discussed in Section 2.1 to get either edge- or coherence-enhancing diffusion. Finally, the 3×3 version G, embedded in $\mathbb{R}^3$ to be used

with PDE (12a), is obtained by reverting the tensor contraction

$$\mathsf{G} = \mathsf{Q}\tilde{\mathsf{G}}\,\mathsf{Q}^{\mathrm{T}}. \tag{29}$$

Note that the flux $\mathbf{j} = \mathsf{G}\nabla_{\mathcal{S}} u$ will automatically be tangential to $\mathcal{S}$ since the columns of Q are a basis of the tangent space.

In (19) we based the definition $\mathsf{J}_{\sigma,\rho}[u](x)$ on integration with respect to surface heat kernels. Instead, in practice, we solve the surface-intrinsic Gaussian diffusion equation

$$\partial_\tau w = \Delta_{\mathcal{S}} w \qquad \text{in } \mathcal{S}, \tau > 0, \tag{30a}$$
$$w|_{\tau=0} = w^0, \tag{30b}$$
$$\nabla_{\mathcal{S}} w^{\mathrm{T}} \boldsymbol{\nu} = 0 \qquad \text{on } \partial\mathcal{S}, \tau > 0\,. \tag{30c}$$

where $\Delta_{\mathcal{S}}$ is the Laplace-Beltrami operator. For the pre-smoothing step we set $w^0 = u$ and solve until $\tau = \sigma$, for the post-smoothing step we initialize w^0 with the ij-component of $\mathsf{J}_{\sigma,0}[u]$ and solve until $\tau = \rho$. Since $\mathsf{J}_{\sigma,\rho}[u]$ is symmetric, this means six further heat solves.

Algorithm 1 summarizes all the steps to get the diffusion tensor G.

Algorithm 1 Construction of $\mathsf{G}[u]$

1: Obtain Q as described in Section 3.2. ▷ (Step 1 can be done once prior to evolving PDE (12a))
2: Find u_σ with a heat solve until σ.
3: Calculate $\nabla_{\mathcal{S}} u_\sigma$ and initialize $\mathsf{J}_{\sigma,0}[u] = \nabla_{\mathcal{S}} u_\sigma \nabla_{\mathcal{S}} u_\sigma^{\mathrm{T}}$
4: Find $\mathsf{J}_{\sigma,\rho}[u]$ with component-wise heat solves until ρ
5: Calculate $\tilde{\mathsf{J}}_{\rho,\sigma}[u] = \mathsf{Q}^{\mathrm{T}} \mathsf{J}_{\sigma,\rho}[u] \mathsf{Q}$
6: Find the spectral decomposition of $\tilde{\mathsf{J}}_{\rho,\sigma}[u]$
7: Set up $\tilde{\mathsf{G}}[u]$ according to edge-enhancing or coherence-enhancing diffusion
8: Calculate $\mathsf{G}[u] = \mathsf{Q}\tilde{\mathsf{G}}[u]\,\mathsf{Q}^{\mathrm{T}}$

3.2 Finding the Tangent Space Basis

We can offer two point-wise approaches to finding an orthonormal basis of the tangent space $\mathcal{T}_x\mathcal{S}$ giving the columns of $\mathsf{Q}(x) = [\mathbf{q}_1(x)|\mathbf{q}_2(x)]$. If $\mathbf{n}$ is given as part of the problem description, then for each point x on the surface, we perform a QR-decomposition (here a single Householder reflection) of the vector $\mathbf{n}(x)$:

$$\mathbf{n}(x) = [\pm\mathbf{n}(x)|\mathbf{q}_1(x)|\mathbf{q}_2(x)]\ (\pm 1, 0, 0)^{\mathrm{T}}. \tag{31}$$

The second and third column are the desired orthonormal tangent vectors.

In Section 4 we will introduce the closest point method which uses a closest point function cp to represent the surface. In [28] we showed that the Jacobian $D\,\mathrm{cp}(x)$ when evaluated at surface points x yields the orthogonal projection matrix $\mathsf{P}(x)$ which projects onto the tangent space $\mathcal{T}_x\mathcal{S}$, i.e.,

$$D\,\mathrm{cp}(x) = \mathsf{P}(x) = \mathsf{I} - \mathbf{n}(x)\mathbf{n}(x)^T \qquad x \in \mathcal{S}. \tag{32}$$

Because of (32) we can find $\mathbf{q}_1(x)$ and $\mathbf{q}_2(x)$ from the spectral decomposition of $D\,\mathrm{cp}(x)$:

$$D\,\mathrm{cp}(x) = [\mathbf{n}|\mathbf{q}_1|\mathbf{q}_2]\,(x) \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} [\mathbf{n}|\mathbf{q}_1|\mathbf{q}_2]\,(x)^T. \tag{33}$$

Now, we have all ingredients of the surface-intrinsic diffusion model. It involves several linear and non-linear in-surface diffusion steps. In the following section we explain how we deal with these equations in order to solve them numerically.

4 Solving In-Surface Diffusion Equations with the Closest Point Method

The closest point method, introduced in [33], is an embedding technique for solving PDEs posed on embedded surfaces. The central idea is to extend functions and differential operators to the surrounding space (here $\mathbb{R}^3$) and to solve an embedding equation which is a 3D-analog of the original surface equation. This approach is appealing since the extended versions of the operators $\nabla_{\mathcal{S}}$, $\mathrm{div}_{\mathcal{S}}$, and $\Delta_{\mathcal{S}}$ in the closest point framework turn out to be ∇, div, and Δ. Hence, embedding equations are accessible to standard finite difference techniques and existing algorithms (and even codes) can be reused.

4.1 Surface Representation

The closest point method utilizes the *closest point representation* of a surface which is given in terms of the closest point function

$$\mathrm{cp}(x) = \arg\min_{\hat{x}\in\mathcal{S}} |x - \hat{x}|. \tag{34}$$

For a point x in the embedding space, $\mathrm{cp}(x)$ is the point on the surface $\mathcal{S}$ which is closest in Euclidean distance to x. This function is well-defined in a tubular neighborhood or narrow band $B(\mathcal{S})$ of the surface and is as smooth as the underlying surface [28].

The closest point function can be derived analytically for many common surfaces. When a parameterization or triangulation of the surface is known, the closest point function can be computed numerically [33]. Figure 3 illustrates the closest point representation of a curve embedded in $\mathbb{R}^2$.

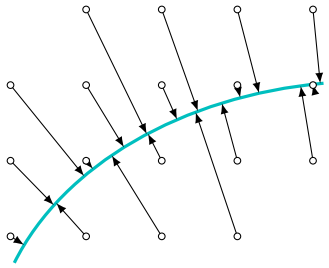

Fig. 3: Closest point representation of a curve embedded in $\mathbb{R}^2$. For each point x (black circles) in the embedding space, the point (arrow tip) on the cyan curve, which is closest in Euclidean distance to x, is stored.

4.2 Extending Functions and Differential Operators

Using the closest point representation, we can extend values of a surface function $u : \mathcal{S} \to \mathbb{R}$ into the surrounding band $B(\mathcal{S})$ by defining $\bar{u} : B(\mathcal{S}) \to \mathbb{R}$ as

$$\bar{u}(x) := u(\mathrm{cp}(x)). \tag{35}$$

Notably, $\bar{u}$ will be constant in the direction normal to the surface and this property is key to the closest point method: it implies that an application of a Cartesian differential operator to $\bar{u}$ is equivalent to applying the corresponding intrinsic surface differential operator to u. We state this as principles below; these mathematical principles were established in [33] and proven in [28].

Principle 1 (Gradient Principle): *Let $\mathcal{S}$ be a surface embedded in $\mathbb{R}^n$ and let u be a function, defined on $\mathbb{R}^n$, that is constant along directions normal to the surface, then*

$$\nabla u(x) = \nabla_{\mathcal{S}} u(x) \qquad \forall\, x \in \mathcal{S}. \tag{36}$$

Principle 2 (Divergence Principle): *Let $\mathcal{S}$ be a surface embedded in $\mathbb{R}^n$. If $\mathbf{j}$ is a vector field on $\mathbb{R}^n$ that is tangent to $\mathcal{S}$ and tangent to all surfaces displaced a fixed Euclidean distance from $\mathcal{S}$, then*

$$\mathrm{div}\, \mathbf{j}(x) = \mathrm{div}_{\mathcal{S}}\, \mathbf{j}(x) \qquad \forall\, x \in \mathcal{S}. \tag{37}$$

Direct consequences of these principles are

$$\nabla_{\mathcal{S}} u(x) = \nabla \bar{u}(x), \tag{38}$$

$$\mathrm{div}_{\mathcal{S}}\, \mathbf{j}(x) = \mathrm{div}\, \bar{\mathbf{j}}(x), \qquad \forall\, x \in \mathcal{S}, \tag{39}$$

where $\bar{\mathbf{j}}$ is the closest point extension of a tangential flux $\mathbf{j}$. Moreover, since $\nabla \bar{u}$ is tangential to level-surfaces of the Euclidean distance-to-$\mathcal{S}$ map [33,28], combining Principles 1 and 2 yields

$$\Delta_{\mathcal{S}} u(x) = \Delta \bar{u}(x), \tag{40}$$

$$\mathrm{div}_{\mathcal{S}}\left(g \nabla_{\mathcal{S}} u\right)(x) = \mathrm{div}\left(\bar{g} \nabla \bar{u}\right)(x), \qquad \forall\, x \in \mathcal{S}, \tag{41}$$

where $g : \mathcal{S} \to \mathbb{R}$ is a scalar diffusivity and $\bar{g}$ its closest point extension.

Additionally, if G is a diffusion tensor which maps to the tangent space, i.e., the corresponding flux $\mathbf{j} = \mathsf{G} \nabla_{\mathcal{S}} u$ is tangential, then Principles 1 and 2 also imply that

$$\mathrm{div}_{\mathcal{S}}\left(\mathsf{G} \nabla_{\mathcal{S}} u\right)(x) = \mathrm{div}\left(\bar{\mathsf{G}} \nabla \bar{u}\right)(x), \qquad \forall\, x \in \mathcal{S}, \tag{42}$$

where $\bar{\mathsf{G}}$ is the closest point extension of G.

Finally, we note that a closest point extension $\bar{u}$ is characterized [18] by

$$\bar{u} = \bar{u} \circ \mathrm{cp}, \tag{43}$$

i.e., it is the closest point extension of itself.

4.3 Gaussian Diffusion

We start with the Gaussian diffusion equation on a closed surface $\mathcal{S}$ in order to demonstrate the embedding idea:

$$\partial_t u = \Delta_{\mathcal{S}} u \qquad x \in \mathcal{S},\ t > 0, \tag{44a}$$

$$u|_{t=0} = u^0. \tag{44b}$$

Using (40) we obtain the following embedding problem

$$\partial_t v = \Delta v \qquad x \in B(\mathcal{S}),\ t > 0, \tag{45a}$$

$$v|_{t=0} = u^0 \circ \mathrm{cp} \tag{45b}$$

$$v = v \circ \mathrm{cp} \qquad x \in B(\mathcal{S}),\ t > 0,\ . \tag{45c}$$

Here (45b) says that we start the process with a closest point extension of the initial data u^0, while condition (45c) guarantees that v stays a closest point extension for all times and hence we can rely on the principles which give us (45a) as the 3D-analog of (44a).

In order to cope with condition (45c) Ruuth & Merriman [33] suggested the following semi-discrete (in time) iteration: after initialization $v^0 = u^0 \circ \mathrm{cp}$, alternate between two steps

1. Evolve $\qquad w = v^n + \tau \Delta v^n,$ (46a)

2. Extend $\qquad v^{n+1} = w \circ \mathrm{cp},$ (46b)

where τ is the time-step size. Here step 1 evolves (45a) of the embedding problem, while step 2 reconstructs the surface function or rather its closest point extension to make sure that (45c) is satisfied at time t_{n+1}.

A fully discrete version of the Ruuth & Merriman iteration needs a discretization of the spatial operators in step 1 and an interpolation scheme in step 2. Using a uniform Cartesian grid in $\mathbb{R}^3$ the Laplacian in our example can be discretized with the standard 7 point

finite difference stencil and be implemented as a matrix L acting on the 1D array v^{n+1}. The interpolation scheme in step 2 is necessary since the data w is given only on grid points and $\mathrm{cp}(x)$ is hardly ever a grid point (even though x is). In order to get around that we interpolate the data w with tri-cubic interpolation and extend the interpolant W rather than w:

$$v_i^{n+1} = W \circ \mathrm{cp}(x_i) \tag{47}$$

where x_i is the grid point corresponding to the i-th component of the array v^{n+1}. Since tri-cubic interpolation is linear in the data w, this operation can also be implemented as an extension matrix E acting on the 1D array w. The fully discrete Ruuth & Merriman iteration reads then

$$\text{1. Evolve} \qquad w = v^n + \tau L v^n, \tag{48a}$$
$$\text{2. Extend} \qquad v^{n+1} = E w. \tag{48b}$$

The computation is performed in a *computational band* for two reasons: firstly, the closest point function is defined in the band $B(\mathcal{S})$ and hence we need sufficient resolution within $B(\mathcal{S})$ to resolve the geometry of the surface. Secondly, the code can be made more efficient by working on a narrow band surrounding the surface $\mathcal{S}$. A nice feature of the Ruuth & Merriman iteration is that no artificial boundary conditions on the boundary of the band need to be imposed [33]. This has to do with the extension step: values at grid points are overwritten at each time step with the value at their closest points. Note also, that no artificial boundary conditions are imposed in the embedding problem (45).

For the sake of efficiency optimization of the width of the computational band is reasonable. The bandwidth depends on the degree of the interpolant and on the finite difference stencil used. Suppose we use an interpolant of degree p and are working in d-dimensions, this requires $(p+1)^d$ points around an interpolation point $\mathrm{cp}(x_i)$ [33]. Furthermore, each of the points in the interpolation stencil must be advanced in time with a finite difference stencil. As a rule of thumb the diameter of the convolution of the interpolation and finite difference stencil gives a good value for the bandwidth. More details on finding the optimal band are given in [26, Appendix A].

4.4 Anisotropic Diffusion

In the case of anisotropic diffusion the closest point embedding idea yields the following embedding problem

$$\partial_t v = \mathrm{div}\left(\mathsf{G}[v] \circ \mathrm{cp}\, \nabla v\right) \qquad x \in B(\mathcal{S}),\ t > 0, \tag{49a}$$
$$v|_{t=0} = u^0 \circ \mathrm{cp} \tag{49b}$$
$$v = v \circ \mathrm{cp} \qquad x \in B(\mathcal{S}),\ t > 0, \tag{49c}$$

which is again a 3D-analog of the surface PDE. Analogously to Gaussian diffusion, we obtain a semi-discrete Ruuth & Merriman iteration, i.e., after initialization $v^0 = u^0 \circ \mathrm{cp}$, we alternate between two steps

$$\text{1. Evolve} \qquad w = v^n + \tau\, \mathrm{div}\left(\mathsf{G}[v^n] \circ \mathrm{cp}\, \nabla v^n\right), \tag{50a}$$
$$\text{2. Extend} \qquad v^{n+1} = w \circ \mathrm{cp}\,. \tag{50b}$$

From here we obtain the discrete iteration: using Algorithm 1 on the state v^n (where all surface-intrinsic differentials and heat solves are dealt with the closest point method) yields the diffusion tensor $\mathsf{G}[v^n] \circ \mathrm{cp}$. The extension (step 2) is realized as explained in the previous section. Finally, we discretize the diffusion operator in (50a) with the following formally second-order accurate scheme:

$$\begin{aligned}
&\mathrm{div}\left(\mathsf{G}[v^n] \circ \mathrm{cp}\, \nabla v^n\right) \approx \\
&\quad D_x^-\left(A_x^+ \mathsf{G}_{11} \mathbin{.*} D_x^+ v^n\right) + D_x^c\left(\mathsf{G}_{12} \mathbin{.*} D_y^c v^n\right) \\
&\quad + D_x^c\left(\mathsf{G}_{13} \mathbin{.*} D_z^c v^n\right) + D_y^c\left(\mathsf{G}_{12} \mathbin{.*} D_x^c v^n\right) \\
&\quad + D_y^-\left(A_y^+ \mathsf{G}_{22} \mathbin{.*} D_y^+ v^n\right) + D_y^c\left(\mathsf{G}_{23} \mathbin{.*} D_z^c v^n\right) \\
&\quad + D_z^c\left(\mathsf{G}_{13} \mathbin{.*} D_x^c v^n\right) + D_z^c\left(\mathsf{G}_{23} \mathbin{.*} D_y^c v^n\right) \\
&\quad + D_z^-\left(A_z^+ \mathsf{G}_{33} \mathbin{.*} D_z^+ v^n\right).
\end{aligned} \tag{51}$$

In this scheme, D_x^-, D_x^+, and D_x^c denote the forward, backward, and central finite difference operators along direction x while A_x^+ averages along direction x to provide values of a diffusion tensor component (in this case G_{11}) on edge-centers on a uniform Cartesian grid. The operator $.*$ means the component wise multiplication of arrays. In the case of isotropic Perona–Malik diffusion, i.e., $\mathsf{G} = g\mathsf{I}$, scheme (51) will reduce to the one we used in [7].

4.5 Non-Dimensionalization and Parameter Adaption

In 2D image processing one typically works on a pixel domain of the form $[0, N-1] \times [0, M-1]$ and a mesh width of $h = 1$. Because our surfaces are contained in a reference box of typical size one, we work with mesh widths $h \ll 1$. Thus, we are in a different scaling regime. In order to compare parameter choices, we non-dimensionalize the diffusion PDEs in a standard fashion (e.g., [20]) and choose parameters relative to the data.

For the linear Gaussian diffusion this results in choosing a new time-scale. To see this, let u satisfy

$$\partial_t u = \Delta u, \qquad x \in [0, L_1] \times [0, L_2] \times [0, L_3],$$

and we look at the transformed solution

$$w(\tau, \xi) = a + b u(\beta\tau, c + L\xi), \qquad L = \max_{i=1,2,3}\{L_i\}, \tag{52}$$

where we have taken into account affine linear transformations of both the color space and the domain. The new function w will then satisfy $\partial_\tau w = \frac{\beta}{L^2}\Delta_\xi w$, and by picking the time scale $\beta = L^2$, we can then solve $\partial_\tau w = \Delta_\xi w$, that is, the same equation for w as we had for u, but to a different stop time. If we are interested in u at time T, we have to evolve w until stop time T/L^2.

In order to get the same scaling behavior for the non-linear Perona–Malik and Weickert models, we adapt the parameters to the initial data (under the assumption that the initial data is non-constant). Starting with the Perona–Malik model

$$\partial_t u = \operatorname{div}\left(g(|\nabla u|^2)\nabla u\right), \qquad g(s^2) = \frac{1}{1+s^2/\lambda^2},$$

we non-dimensionalize the nonlinear diffusivity g by taking into account the initial data u^0 and set

$$\lambda = \lambda_{\mathrm{rel}} \|\nabla u^0_\sigma\|_\infty.$$

That is, a gradient is considered to be large if its magnitude is above a certain percentage λ_{rel} of the gradient magnitude of the given data u^0 (or precisely, the linearly smoothed version u^0_σ). Let w be again the transformation of (52). Taking into account the appropriate time-scale for the linear Gaussian pre-smoothing, the data transforms as

$$w^0(\xi) = a + bu^0(L\xi), \qquad w^0_{\frac{\sigma}{L^2}}(\xi) = a + bu^0_\sigma(L\xi).$$

Consequently, the ratio

$$\frac{|\nabla_\xi w|}{\|\nabla_\xi w^0_{\frac{\sigma}{L^2}}\|_\infty} = \frac{|\nabla u|}{\|\nabla u^0_\sigma\|_\infty}$$

is invariant and hence we obtain

$$\partial_\tau w = \frac{\beta}{L^2}\operatorname{div}_\xi\left(g(|\nabla_\xi w|^2)\nabla_\xi w\right).$$

By using the time-scale $\beta = L^2$ (as was the case for Gaussian diffusion) we end up with the original Perona–Malik model.

In Weickert's anisotropic diffusion models the nonlinear functions depend on the coherence, so we choose the parameters relative to the coherence of the initial data. For edge-enhancing and coherence-enhancing diffusion the eigenvalues of the diffusion tensor G are given by (7) and (9) respectively, in terms of the coherence c. Let c^0 denote the coherence of the initial data, we set

$$\lambda = \lambda_{\mathrm{rel}} \|c^0\|_\infty, \qquad B = B_{\mathrm{rel}} \|c^0\|_\infty. \tag{53}$$

Because of the scaling behavior of Gaussian diffusion we have

$$\mathsf{J}_{\sigma,\rho}[u] = \frac{1}{b^2L^2}\mathsf{J}_{\frac{\sigma}{L^2},\frac{\rho}{L^2}}[w].$$

Thus, with appropriate time-scales for the Gaussian pre-smoothing and post-smoothing steps, the ratio $c/\|c^0\|_\infty$ is invariant under affine linear transformations. We obtain the transformed PDE

$$\partial_\tau w = \frac{\beta}{L^2}\operatorname{div}_\xi\left(\mathsf{G}[w]\nabla_\xi w\right).$$

The time scale factor is thus again $\beta = L^2$.

By virtue of (13), the surface-intrinsic diffusion PDEs have the same scaling behavior as their $\mathbb{R}^3$-counterparts.

5 Experiments

We demonstrate our algorithm on several examples. The numerical parameters used are the mesh width $h = 0.0125$ and the time-step size $\tau = 0.15h^2$. We use a uniform Cartesian grid defined on the reference box $[-1.5, 1.5]^3$, but the algorithm is executed on a narrow computational band containing only grid points close to the surface.

5.1 Denoising

In Figures 4 and 5 we compare Gaussian, Perona–Malik, and edge-enhancing diffusion.

5.1.1 Stripe Pattern on a Torus

The top left image of Figures 4 shows a noisy stripe pattern defined on a torus. The torus is the solution of

$$(R - \sqrt{x^2+y^2})^2 + z^2 = r^2 \tag{54}$$

where the big and the small radius are $R = 1$ and $r = 0.4$. From (54) we find the normal field analytically and obtain the tangent space basis, which is needed in edge-enhancing diffusion, by QR-decomposition of the normal.

In Figure 4 we observe the usual effect of Gaussian diffusion: edges are not sharp because of the uniform blur. Perona–Malik diffusion produces sharp edges, but the noise on the edge is still present since the flux vanishes on edges. In contrast to that, edge-enhancing diffusion aligns the flux with edges and produces sharp but smoother edges.

5.1.2 Wood Grain on a Sphere

The top left image of Figures 5 shows a wood grain defined on the unit sphere. The normal field is found analytically from the defining equation $x^2+y^2+z^2 = 1$ and the tangent space basis is obtained by QR-decomposition of the normal.

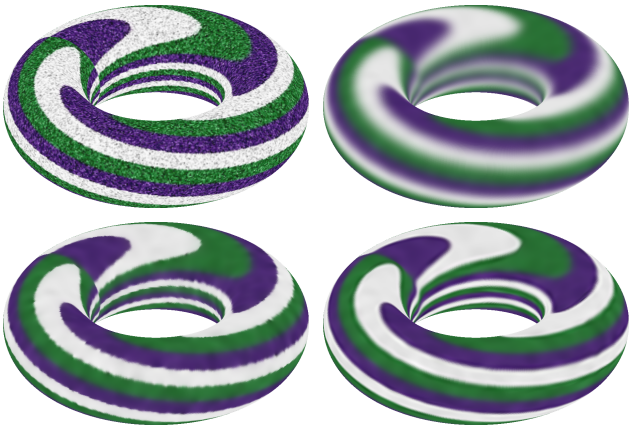

Fig. 4: Denoising on a torus: the effect of different diffusion filters. *Top Left:* Noisy image. *Top Right:* Gaussian diffusion. *Bottom Left:* Perona–Malik diffusion; edge-sensitivity parameter $\lambda_{\text{rel}} = 2 \cdot 10^{-1}$. *Bottom Right:* Edge-enhancing diffusion; $\sigma = 1 \cdot 10^{-4}$, $\rho = 4 \cdot 10^{-4}$, $\lambda_{\text{rel}} = 4 \cdot 10^{-2}$. Stop time in all cases $T = 1.2 \cdot 10^{-3}$ (52 iterations).

In Figure 5 the effects are even more drastic than in Figure 4: Gaussian diffusion removes the noise but does hardly preserve any structure. Perona–Malik diffusion preserves some structures, but most of the noise is still present. In contrast to that, edge-enhancing diffusion removes the noise while preserving the structure of the wood grain.

Fig. 5: Denoising on a sphere: the effect of different diffusion filters. *Top Left:* Noisy image. *Top Right:* Gaussian diffusion. *Bottom Left:* Perona–Malik diffusion; edge-sensitivity parameter $\lambda_{\text{rel}} = 2 \cdot 10^{-1}$. *Bottom Right:* Edge-enhancing diffusion; $\sigma = 1 \cdot 10^{-4}$, $\rho = 4 \cdot 10^{-4}$, $\lambda_{\text{rel}} = 4 \cdot 10^{-2}$. Stop time in all cases $T = 5.9 \cdot 10^{-4}$ (25 iterations).

5.2 Coherence Enhancement

5.2.1 Fingerprint on a Sphere

The fingerprint image is taken from Weickert [38], here we texture-mapped it to the lower and upper hemisphere of a unit sphere (perhaps with the upper sphere serving as simple model of finger tip). Since the unit sphere is the solution of $x^2 + y^2 + z^2 = 1$ we find the normal field analytically and obtain the tangent space basis by QR-decomposition of the normal. Our result, bottom right image of Figure 6, is similar to Weickert's result, top right image of Figure 6. In the following, we compare parameter choices.

Weickert reports the following parameters for his experiment with the fingerprint (an image of size 256×256 pixels with values in $[0, 255]$): $\alpha = 10^{-3}$, $B = 1$ for the diffusion tensor and $T = 20$ as the stop time for the anisotropic diffusion. Regarding the structure tensor $\sigma = 0.5$, $\rho = 4$ were used.

Our upper hemisphere has surface area 2π so a reasonable length factor based on area is $L = 255/\sqrt{2\pi} \approx 10^2$. In the set-up of the structure tensor, Weickert convolves with Gaussian kernels while we are (formally) convolving with heat-kernels. Weickert's time scales would transform to

$$\hat{\sigma} = \frac{\sigma^2}{2L^2} \approx 1.2 \cdot 10^{-5}, \qquad \hat{\rho} = \frac{\rho^2}{2L^2} \approx 7.7 \cdot 10^{-4},$$
$$\hat{T} = \frac{T}{L^2} \approx 1.9 \cdot 10^{-3}.$$

Finally we take into account our relative choice of B as discussed above. The maximal coherence of the initial state (top left image of Figure 6) is about $\|c^0\|_\infty \approx 10^3$. Thus, in order to obtain $B = 1$, we have $B_{\text{rel}} \approx 10^{-3}$.

The parameters which we actually used in our algorithm are

$$\sigma = 10^{-4}, \qquad \rho = 4 \cdot 10^{-4}, \qquad T = 1.2 \cdot 10^{-3},$$
$$\alpha = 10^{-3}, \qquad B_{\text{rel}} = 10^{-3}, \tag{55}$$

and are comparable (in order of magnitude) to Weickert's.

5.2.2 Sunflowers on a Vase

This example demonstrates the stylization of images as an application of coherence enhancing diffusion. The

Fig. 6: Feature enhancement in a fingerprint. *Top Left:* Fingerprint image (by courtesy of J. Weickert). *Top Right:* Weickert's result with coherence enhancing diffusion [38]. *Bottom Left:* Fingerprint texture-mapped to a sphere. *Bottom Right:* Our result with in-surface coherence enhancing diffusion with the parameter values of (55).

Fig. 7: Stylization of a sunflower picture on a vase. *Top Left:* Original image. Results of coherence enhancing diffusion with the parameter values of (56): *Top Right:* iteration 10 ($T = 2.3 \cdot 10^{-4}$), *Bottom Left:* iteration 20 ($T = 4.7 \cdot 10^{-4}$), *Bottom Right:* iteration 30 ($T = 7.0 \cdot 10^{-4}$).

vase is a surface of revolution with a sunflower texture map. Figure 7 shows the results obtained with the parameters

$$\sigma = 10^{-4}, \quad \rho = 5 \cdot 10^{-4}, \quad \alpha = 10^{-3}, \quad B_{\text{rel}} = 10^{-3}. \quad (56)$$

Since this an open surface, we must impose a boundary condition on the edge of the vase. One reasonable choice would be the no-flux boundary condition (12c): $\boldsymbol{\nu}^T \mathbf{j} = 0$ at the boundary of $\mathcal{S}$. Another possible choice is the Neumann-zero condition $\boldsymbol{\nu}^T \nabla_{\mathcal{S}} u = 0$, which we use here because it is easy to implement with the closest point method (as it is automatically satisfied in the closest point framework). However, it is likely not conservative. Nonetheless, the results of Figure 7 are visually reasonable at the edges.

5.2.3 Impressionist Style on a Triangulated Surface

As another example of stylization of images using coherence enhancing diffusion, we consider texture creation. Here we have a triangulation of an ear, originally from [32] and smoothed and Loop subdivided to a smooth mesh of roughly 400 000 triangles. We convert the triangulation to a closest point function using the technique outlined in [25]. The basis for the tangent space is computed from the closest point function using the derivative of the projection as in (33). In Figure 8 (top-left), an artist (ahem) has quickly and roughly painted some parts of the surface using colors chosen from the palette of van Gogh's "Starry Night". We perturb this input image by setting randomly 67% of the voxels to a random selection of mostly dark blues, again from the palette of "Starry Night". This noisy image is shown in Figure 8 (top-right). The ear is two units long in its longest dimension and we embed it in a grid with $h = 0.008$ (smaller than the other examples to get a finer texture). The other parameters are chosen as

$$\sigma = 10^{-4}, \quad \rho = 5 \cdot 10^{-4}, \quad \alpha = 10^{-3}, \quad B_{\text{rel}} = 10^{-6}.$$

Figure 8 (bottom) shows the results, a simulated Impressionist painting on the surface of an ear.

6 Conclusions

In this paper we introduced a model for edge- and coherence-enhancing image processing on curved surfaces using surface-intrinsic anisotropic diffusion. We defined a surface-intrinsic structure tensor, from which

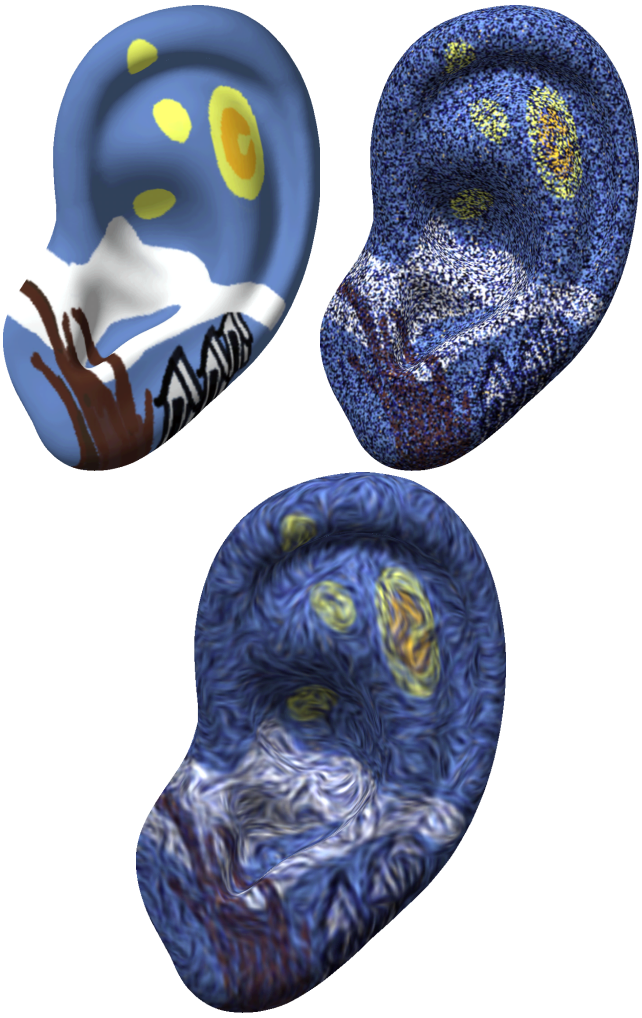

Fig. 8: "Starry Night" on an ear: creating Impressionist textures. *Top-left:* user input: some colors roughly painted onto the surface, *Top-right:* 67% of pixels replaced with dark blue shades, *Bottom:* Coherence enhancement, iteration 10 ($T = 1.1 \cdot 10^{-4}$).

the construction of the diffusion tensor follows almost exactly the procedure suggested by Weickert. The resulting surface-intrinsic diffusion PDE is solved numerically using the closest point method, a general method for solving PDEs posed on surfaces. Our approach can be used for denoising of data posed on surfaces and for visual effects such as generating surface textures or stylizing existing textures. Our results for images on surfaces are comparable to those of Weickert and parameter choices made for 2D images can be re-used taking into account the corresponding scale factors.

For open surfaces we used zero-Neumann rather than no-flux boundary conditions. The implementation of general no-flux boundary conditions within the closest point method is a topic for future research.

Acknowledgements This work was supported by award KUK-C1-013-04 made by King Abdullah University of Science and Technology (KAUST).

References

1. T. Aach, C. Mota, I. Stuke, M. Mühlich, and E. Barth. Analysis of superimposed oriented patterns. *IEEE Transactions on Image Processing*, 15(12):3690–3700, 2006.
2. L. Alvarez, P. Lions, and J. Morel. Image selective smoothing and edge detection by nonlinear diffusion. II. *SIAM Journal on Numerical Analysis*, 29(3):845–866, 1992.
3. S. Auer, C.B. Macdonald, M. Treib, J. Schneider, and R. Westermann. Real-time fluid effects on surfaces using the Closest Point Method. *Comput. Graph. Forum*, 31(6):1909–1923, 2012.
4. C.L. Bajaj and G. Xu. Anisotropic diffusion of surfaces and functions on surfaces. *ACM Trans. Graphic.*, 22(1):4–32, 2003.
5. M. Bertalmìo, L.-T. Cheng, S. Osher, and G. Sapiro. Variational problems and partial differential equations on implicit surfaces. *J. Comp. Phys.*, 174(2), 2001.
6. M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester. Image inpainting. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 417–424, 2000.
7. H. Biddle, I. von Glehn, C. B. Macdonald, and T. März. A volume-based method for denoising on curved surfaces. In *2013 20th IEEE International Conference on Image Processing (ICIP)*, pages 529–533. IEEE, 2013.
8. I. Bogdanova, X. Bresson, J.-P. Thiran, and P. Vandergheynst. Scale-space analysis and active contours for omnidirectional images. *IEEE Trans. Image Process.*, 16(7):1888–1901, 2007.
9. R.G. Brown, C.E. Hann, and J.G. Chase. Vision-based 3D surface motion capture for the DIET breast cancer screening system. *Int. J. Comput. Appl. T.*, 39(1):72–78, 2010.
10. T. Brox, J. Weickert, B. Burgeth, and P. Mrázek. Nonlinear structure tensors. *Image and Vision Computing*, 24(1):41–55, 2006.
11. F. Catté, P. Lions, J. Morel, and T. Coll. Image selective smoothing and edge detection by nonlinear diffusion. *SIAM Journal on Numerical Analysis*, 29(1):182–193, 1992.
12. T.F. Chan and J. Shen. Mathematical models for local nontexture inpaintings. *SIAM Journal on Applied Mathematics*, 62(3):1019–1043, 2002.
13. T.F. Chan and L.A. Vese. Active contours without edges. *IEEE Transactions on Image Processing*, 10(2):266–277, 2001.
14. U. Clarenz, U. Diewald, and M. Rumpf. Processing textured surfaces via anisotropic geometric diffusion. *IEEE Trans. Image Process.*, 13(2):248–261, 2004.
15. M. Desbrun, M. Meyer, P. Schröder, and A.H. Barr. Anisotropic feature-preserving denoising of height fields and bivariate data. In *Graphics Interface*, volume 11, pages 145–152, 2000.
16. G. Dziuk and C.M. Elliott. Surface finite elements for parabolic equations. *J. Comp. Math.*, 25(4):385–407, 2007.
17. J. Faraudo. Diffusion equation on curved surfaces. I. Theory and application to biological membranes. *Journal of Chemical Physics*, 116(13):5831–5841, 2002.

18. I. von Glehn, T. März, and C.B. Macdonald. An embedded method-of-lines approach to solving PDEs on surfaces. submitted, 2013.
19. K. Hara, R. Kurazume, K. Inoue, and K. Urahama. Segmentation of images on polar coordinate meshes. In *2007 IEEE International Conference on Image Processing*, volume 2, pages II–245–II–248, 2007.
20. M. H. Holmes. *Introduction to the Foundations of Applied Mathematics*. Springer, 2009.
21. R. Kimmel. Intrinsic scale space for images on surfaces: The geodesic curvature flow. In *Graphical Models and Image Processing*, pages 365–372. Springer-Verlag, 1997.
22. J.J. Koenderink. The structure of images. *Biological Cybernetics*, 50(5):363–370, 1984.
23. R.-J. Lai and T.F. Chan. A framework for intrinsic image processing on surfaces. *Comput. Vis. Image Und.*, 115(12):1647–1661, 2011.
24. Y. Lin and W.B. Seales. Opaque document imaging: Building images of inaccessible texts. In *Proc. ICCV'05, 10th IEEE International Conference on Computer Vision vol. 1*, 2005.
25. C.B. Macdonald and S.J. Ruuth. Level set equations on surfaces via the Closest Point Method. *J. Sci. Comput.*, 35(2–3), 2008.
26. C.B. Macdonald and S.J. Ruuth. The implicit closest point method for the numerical solution of partial differential equations on surfaces. *SIAM Journal on Scientific Computing*, 31(6):4330–4350, 2010.
27. T. März. Image inpainting based on coherence transport with adapted distance functions. *SIAM Journal on Imaging Sciences*, 4(4):981–1000, 2011.
28. T. März and C.B. Macdonald. Calculus on surfaces with general closest point functions. *SIAM J. Numer. Anal.*, 50(6):3303–3328, 2012.
29. J.D. Murray. A pre-pattern formation mechanism for animal coat markings. *Journal of Theoretical Biology*, 88(1):161–199, 1981.
30. I.L. Novak, F. Gao, Y. Choi, D. Resasco, J.C. Schaff, and B.M. Slepchenko. Diffusion on a curved surface coupled to diffusion in the volume: Application to cell biology. *Journal of Computational Physics*, 226(2):1271–1290, 2007.
31. P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7):629–639, 1990.
32. H. Rusinoff. 3D model of human ear. http://www.turbosquid.com/FullPreview/Index.cfm/ID/483517, 2009. Accessed 2014-03-10.
33. S.J. Ruuth and B. Merriman. A simple embedding method for solving partial differential equations on surfaces. *J. Comput. Phys.*, 227(3), 2008.
34. E. Sander and T. Wanner. Pattern formation in a nonlinear model for animal coats. *Journal of Differential Equations*, 191(1):143–174, 2003.
35. P. Schwartz, D. Adalsteinsson, P. Colella, A.P. Arkin, and M. Onsum. Numerical computation of diffusion on a surface. *Proceedings of the National Academy of Sciences of the United States of America*, 102(32):11151–11156, 2005.
36. A. Spira and R. Kimmel. Geometric curve flows on parametric manifolds. *J. Comput. Phys.*, 223, 2007.
37. L. Tian, C.B. Macdonald, and S.J. Ruuth. Segmentation on surfaces with the Closest Point Method. In *Proc. ICIP09, 16th IEEE International Conference on Image Processing*, pages 3009–3012, 2009.
38. J. Weickert. *Anisotropic Diffusion in Image Processing*. B.G. Teubner Stuttgart, 1998.
39. J. Weickert. Coherence-enhancing diffusion of colour images. *Image and Vision Computing*, 17(3):201–212, 1999.
40. Y. Zhai, D. Zhang, J. Sun, and B. Wu. A novel variational model for image segmentation. *Journal of Computational and Applied Mathematics*, 235(8):2234–2241, 2011.